

WHAT IS CLAIMED IS:

1 1. A method for sending data, comprising:
2 establishing an active connection adapted to send packets of data between a host and a
3 destination;
4 receiving from the destination a first window value representing a first quantity of data
5 packets;
6 sending packets of data from said host to said destination;
7 receiving an acknowledgment from said destination for each packet of data received by
8 said destination wherein said first window value represents a limit imposed on said host by said
9 destination on the quantity of data packets sent from said host to said destination and lacking an
10 acknowledgment of being received by destination; and
11 limiting the number of packets sent by said host, but not acknowledged as received by
12 said destination, to a second quantity of data packets less than said first window value.

1 2. The method of claim 1 wherein the connection is a Transmission Control
2 Protocol connection between the host and the destination and wherein said first window value
3 is a Transmission Control Protocol send window value.

1 3. The method of claim 1 further comprising:
2 establishing a plurality of active connections between the host and a plurality of
3 destinations;
4 receiving from each destination a first window value representing a first quantity of data
5 packets for the connection;
6 sending packets of data from said host to each destination;
7 receiving an acknowledgment from each destination for each packet of data received
8 by each destination wherein the first window value of each connection represents a limit
9 imposed on said host by the destination of the connection on the quantity of data packets sent

10 from said host to the destination of the connection and lacking an acknowledgment of being
11 received by the destination of the connection; and
12 limiting the number of packets sent by said host to each connection, but not
13 acknowledged as received by the destination of each connection, to a second quantity of data
14 packets less than the window value of the connection;
15 wherein the second quantity of each connection which is less than the window value of
16 the connection is based, at least in part, on the number of active connections of the host.

1 4. The method of claim 2 wherein said host has a plurality of Transmission Control
2 Protocol connections, each Transmission Control Protocol connection having a Protocol
3 Control Block which stores a Transmission Control Protocol send window value and a virtual
4 window value less than said Transmission Control Protocol send window value wherein each
5 virtual window value limits the number of packets sent by said host, but not acknowledged as
6 received by the destination of each Transmission Control Protocol connection, to a second
7 quantity of data packets defined by the virtual window value of the Transmission Control
8 Protocol connection.

1 5. The method of claim 3 further comprising,
2 in response to the destination reducing the size of the Transmission Control Protocol
3 send window value to a third quantity less than the second quantity, limiting the number of
4 packets sent by said host, but not acknowledged as received by said destination, to a fourth
5 quantity of data packets no greater than the reduced size of the Transmission Control Protocol
6 send window value.

1 6. The method of claim 1, further comprising:
2 establishing a plurality of active direct memory access connections between said host
3 and a plurality of specified memory locations of a plurality of destinations;

4 sending a plurality of messages to specified memory locations of the destinations of the
5 direct memory access connections wherein each message comprises a plurality of data packets;
6 receiving message acknowledgments, each message acknowledgment being sent by a
7 destination for each message received by the destination; and
8 establishing a plurality of message limits, each message limit imposing a separate limit
9 for each direct memory access connection on the quantity of messages sent from said host to
10 the specified memory location of the direct memory access connection associated with the
11 message limit and lacking a message acknowledgment of being received by the destination of
12 the direct memory access connection associated with the message limit.

1 7. The method of claim 6, wherein each direct memory access connection includes
2 a network interface between an application of said host and a network connecting the host to
3 the plurality of destinations and wherein said network interface includes a queue for each direct
4 memory access connection and adapted to queue messages to be sent through the direct
5 memory access connection associated with each queue, and wherein said each sending of a
6 message to specified memory location of the destination of a direct memory access connection
7 includes queuing the message in the network interface queue associated with the direct memory
8 access connection; and wherein the queuing of messages in the network interface queue
9 associated with a direct memory access connection is suspended when the quantity of messages
10 sent from said host to the specified memory location of the associated direct memory access
11 connection and lacking a message acknowledgment of being received by the destination of the
12 associated direct memory access connection reaches the separate message limit imposed on the
13 direct memory access connection associated with the network interface queue.

1 8. The method of claim 7, wherein the queuing of messages in the network
2 interface queue associated with a direct memory access connection is resumed when the
3 quantity of messages sent from said host to the specified memory location of the associated

4 direct memory access connection and lacking a message acknowledgment of being received by
5 the destination of the associated direct memory access connection is less than the separate
6 message limit imposed on the direct memory access connection associated with the network
7 interface queue.

1 9. The method of claim 8 wherein the packet sending connection is a
2 Transmission Control Protocol connection between the host and the destination and wherein
3 each direct memory access connection is a Remote Direct Memory Access connection
4 between the host and the destination of the direct memory access connection.

1 10. The method of claim 10 wherein said network interface has a pool of empty
2 messages which imposes a limit on the total quantity of messages sent from said host to all the
3 specified memory locations of all the direct memory access connections and lacking a message
4 acknowledgment of being received by the destination of the associated direct memory access
5 connection and wherein each message limit is less than the network interface pool of empty
6 messages.

1 11. The method of claim 6 wherein each message limit is based, at least in part, on
2 the number of active direct memory access connections of the host.

1 12. The method of claim 6 further comprising changing the size of a message limit of
2 an active direct memory access connection prior to sending some messages through the
3 associated direct memory access connection.

1 13. The method of claim 6 wherein each message limit is based, at least in part, on
2 the number of active direct memory access connections of the host.

1 14. The method of claim 1 further comprising changing the size of the second
2 quantity of packets limiting the number of packets sent by the host but not acknowledged as
3 received by the destination prior to sending some packets.

1 15. A system adapted to communicate with data storage and a destination having
2 memory locations, comprising:

3 a system memory;
4 a processor coupled to the system memory;
5 a network adaptor;
6 a data storage controller for managing Input/Output (I/O) access to the data storage;
7 and

8 a device driver executable by the processor in the memory, wherein at least one of the
9 device driver and the network adaptor is adapted to:

10 (i) establish an active connection adapted to send packets of data between the
11 system and a destination;

12 (ii) receive from a destination a first window value representing a first quantity
13 of data packets;

14 (iii) send packets of data from the system to said destination;

15 (iv) receive an acknowledgment from said destination for each packet of data
16 received by said destination wherein said first window value represents a limit imposed
17 on said system by said destination on the quantity of data packets sent from said system
18 to said destination and lacking an acknowledgment of being received by destination;
19 and

20 (v) limit the number of packets sent by said system, but not acknowledged as
21 received by said destination, to a second quantity of data packets less than said first
22 quantity.

1 16. The system of claim 15, wherein the data storage comprises a magnetic storage
2 medium.

1 17. The system of claim 15, wherein the connection is a Transmission Control
2 Protocol connection between the host and the destination and wherein said first window value
3 is a Transmission Control Protocol send window value.

1 18. The system of claim 15 wherein at least one of the device driver and the
2 network adaptor is further adapted to:
3 establish a plurality of active connections between the system and a plurality of
4 destinations;
5 receive from each destination a first window value representing a first quantity of data
6 packets for the connection;
7 send packets of data from said system to each destination;
8 receive an acknowledgment from each destination for each packet of data received by
9 each destination wherein the first window value of each connection represents a limit imposed
10 on said system by the destination of the connection on the quantity of data packets sent from
11 said system to the destination of the connection and lacking an acknowledgment of being
12 received by the destination of the connection; and
13 limit the number of packets sent by said system to each connection, but not
14 acknowledged as received by the destination of each connection, to a second quantity of data
15 packets less than the window value of the connection;
16 wherein the second quantity of each connection which is less than the window value of
17 the connection is based, at least in part on the number of active connections of the system.

1 19. The system of claim 17 wherein at least one of the device driver and the
2 network adaptor is adapted to establish a plurality of Transmission Control Protocol
3 connections, each Transmission Control Protocol connection having a Protocol Control Block
4 data structure which stores a Transmission Control Protocol send window value and a virtual
5 window value less than said Transmission Control Protocol send window value wherein each
6 virtual window value limits the number of packets sent by said system, but not acknowledged as
7 received by the destination of each Transmission Control Protocol connections, to a second
8 quantity of data packets defined by the virtual window value of the Transmission Control
9 Protocol connection.

1 20. The system of claim 19 wherein at least one of the device driver and the
2 network adaptor is adapted to:
3 in response to the destination reducing the size of the Transmission Control Protocol
4 send window value to a third quantity less than the second quantity, limiting the number of
5 packets sent by said system, but not acknowledged as received by said destination, to a fourth
6 quantity of data packets no greater than the reduced size of the Transmission Control Protocol
7 send window value.

1 21. The system of claim 15, wherein at least one of the device driver and the
2 network adaptor is adapted to:
3 establish a plurality of active direct memory access connections between said system
4 and a plurality of specified memory locations of a plurality of destinations;
5 send a plurality of messages to specified memory locations of the destinations of the
6 direct memory access connections wherein each message comprises a plurality of data packets;
7 receive message acknowledgments, each message acknowledgment being sent by a
8 destination for each message received by the destination; and

9 establish a plurality of message limits, each message limit imposing a separate limit for
10 each direct memory access connection on the quantity of messages sent from said system to the
11 specified memory location of the direct memory access connection associated with the message
12 limit and lacking a message acknowledgment of being received by the destination of the direct
13 memory access connection associated with the message limit.

1 22. The system of claim 21, wherein at least one of the device driver and the
2 network adaptor is adapted to provide a queue for each direct memory access connection and
3 adapted to queue messages to be sent through the direct memory access connection associated
4 with each queue, and wherein in each sending of a message to specified memory location of the
5 destination of a direct memory access connection, at least one of the device driver and the
6 network adaptor is adapted to queue the message in the queue associated with the direct
7 memory access connection; and to suspend the queuing of messages in the queue associated
8 with a direct memory access connection when the quantity of messages sent from said system
9 to the specified memory location of the associated direct memory access connection and
10 lacking a message acknowledgment of being received by the destination of the associated direct
11 memory access connection reaches the separate message limit imposed on the direct memory
12 access connection associated with the queue.

1 23. The system of claim 22, wherein at least one of the device driver and the
2 network adaptor is adapted to resume the queuing of messages in the queue associated with a
3 direct memory access connection when the quantity of messages sent from said system to the
4 specified memory location of the associated direct memory access connection and lacking a
5 message acknowledgment of being received by the destination of the associated direct memory
6 access connection is less than the separate message limit imposed on the direct memory access
7 connection associated with the network interface queue.

1 24. The system of claim 23 wherein the packet sending connection is a
2 Transmission Control Protocol connection between the system and the destination of the packet
3 sending connection and wherein each direct memory access connection is a Remote Direct
4 Memory Access connection between the system and the destination of the direct memory
5 access connection.

1 25. The system of claim 24 wherein at least one of the device driver and the
2 network adaptor is adapted to provide a pool of empty messages which imposes a limit on the
3 total quantity of messages sent from said system to all the specified memory locations of all the
4 direct memory access connections and lacking a message acknowledgment of being received
5 by the destination of the associated direct memory access connection and wherein each
6 message limit is less than the pool of empty messages.

1 26. An article of manufacture for managing data transmission through a network,
2 wherein the article of manufacture causes operations to be performed, the operations
3 comprising:
4 establishing an active connection adapted to send packets of data between a host and a
5 destination;
6 receiving from a destination a first window value representing a first quantity of data
7 packets;
8 sending packets of data from a host to said destination;
9 receiving an acknowledgment from said destination for each packet of data received by
10 said destination wherein said first window value represents a limit imposed on said host by said
11 destination on the quantity of data packets sent from said host to said destination and lacking an
12 acknowledgment of being received by destination; and
13 limiting the number of packets sent by said host, but not acknowledged as received by
14 said destination, to a second quantity of data packets less than said first quantity.

1 27. The article of manufacture of claim 26 wherein the connection is a Transmission
2 Control Protocol connection between the host and the destination and wherein said first
3 window value is a Transmission Control Protocol send window value.

1 28. The article of manufacture of claim 27 wherein the operations further comprise:
2 establishing a plurality of active connections between the host and a plurality of
3 destinations;
4 receiving from each destination a first window value representing a first quantity of data
5 packets for the connection;
6 sending packets of data from said host to each destination;
7 receiving an acknowledgment from each destination for each packet of data received
8 by each destination wherein the first window value of each connection represents a limit
9 imposed on said host by the destination of the connection on the quantity of data packets sent
10 from said host to the destination of the connection and lacking an acknowledgment of being
11 received by the destination of the connection; and
12 limiting the number of packets sent by said host to each connection, but not
13 acknowledged as received by the destination of each connection, to a second quantity of data
14 packets less than the window value of the connection;
15 wherein the second quantity of each connection which is less than the window value of
16 the connection is based, at least in part, on the number of active connections of the host.

1 29. The article of manufacture of claim 28 wherein said host has a plurality of
2 Transmission Control Protocol connections, each Transmission Control Protocol connection
3 having a Protocol Control Block which stores a Transmission Control Protocol send window
4 value and a virtual window value less than said Transmission Control Protocol send window
5 value wherein each virtual window value limits the number of packets sent by said host, but not

6 acknowledged as received by the destination of each Transmission Control Protocol
7 connections, to a second quantity of data packets defined by the virtual window value of the
8 Transmission Control Protocol connection.

1 30. The article of manufacture of claim 28, wherein the operations further comprise:
2 in response to the destination reducing the size of the Transmission Control Protocol
3 send window value to a third quantity less than the second quantity, limiting the number of
4 packets sent by said host, but not acknowledged as received by said destination, to a fourth
5 quantity of data packets no greater than the reduced size of the Transmission Control Protocol
6 send window value.

1 31. The article of manufacture of claim 26, wherein the operations further comprise:
2 establishing a plurality of active direct memory access connections between said host
3 and a plurality of specified memory locations of a plurality of destinations;
4 sending a plurality of messages to specified memory locations of the destinations of the
5 direct memory access connections wherein each message comprises a plurality of data packets;
6 receiving message acknowledgments, each message acknowledgment being sent by a
7 destination for each message received by the destination;
8 establishing a plurality of message limits, each message limit imposing a separate limit
9 for each direct memory access connection on the quantity of messages sent from said host to
10 the specified memory location of the direct memory access connection associated with the
11 message limit and lacking a message acknowledgment of being received by the destination of
12 the direct memory access connection associated with the message limit.

13

1 32. The article of manufacture of claim 31, wherein each direct memory access
2 connection includes a network interface between an application of said host and a network
3 connecting the host to the plurality of destinations and wherein said network interface includes a

4 queue for each direct memory access connection and adapted to queue messages to be sent
5 through the direct memory access connection associated with each queue, and wherein said
6 each sending of a message to specified memory location of the destination of a direct memory
7 access connection includes queuing the message in the network interface queue associated with
8 the direct memory access connection; and wherein the queuing of messages in the network
9 interface queue associated with a direct memory access connection is suspended when the
10 quantity of messages sent from said host to the specified memory location of the associated
11 direct memory access connection and lacking a message acknowledgment of being received by
12 the destination of the associated direct memory access connection reaches the separate
13 message limit imposed on the direct memory access connection associated with the network
14 interface queue.

1 33. The article of manufacture of claim 32, wherein the queuing of messages in the
2 network interface queue associated with a direct memory access connection is resumed when
3 the quantity of messages sent from said host to the specified memory location of the associated
4 direct memory access connection and lacking a message acknowledgment of being received by
5 the destination of the associated direct memory access connection is less than the separate
6 message limit imposed on the direct memory access connection associated with the network
7 interface queue.

1 34. The article of manufacture of claim 33 wherein the packet sending connection is
2 a Transmission Control Protocol connection between the host and the destination and wherein
3 each direct memory access connection is a Remote Direct Memory Access connection
4 between the host and the destination of the direct memory access connection.

1 35. The article of manufacture of claim 34 wherein said network interface has a
2 pool of empty messages which imposes a limit on the total quantity of messages sent from said

3 host to all the specified memory locations of all the direct memory access connections and
4 lacking a message acknowledgment of being received by the destination of the associated direct
5 memory access connection and wherein each message limit is less than the network interface
pool of empty messages.

1 36. The article of manufacture of claim 31 wherein each message limit is based, at
2 least in part, on the number of active direct memory access connections of the host.

1 37. The article of manufacture of claim 31 further comprising changing the size of a
2 message limit of an active direct memory access connection prior to sending some messages
3 through the associated direct memory access connection.

1 38. The article of manufacture of claim 31 wherein each message limit is based, at
2 least in part, on the number of active direct memory access connections of the host.

1 39. The article of manufacture of claim 26 further comprising changing the size of
2 the second quantity of packets limiting the number of packets sent by the host but not
3 acknowledged as received by the destination prior to sending some packets.